

Thousands of Linux Installations (and only one administrator)

A Linux cluster client
for the University of Manchester

A V Le Blanc
Manchester Computing
University of Manchester

LeBlanc@mcc.ac.uk

Overview

- Description of the system
- How to look after it
- Why is it this way?
- Evaluation

Description of the System

- Base layer
- Files on disk
- Cache
- Full system on the network

Description of the System

Base	4Mb
------	-----

On disk	28Mb
---------	------

Cache	500Mb
-------	-------

Full system	2.3Gb
-------------	-------

Base Layer

- Kernel 918 Kb
- Ramdisk 2825 Kb
- Config file 50 bytes
- In a separate partition, also containing:
 - GRUB
 - Novell's ZenWorks

Base Layer (booting)

- Boot from this kernel with ramdisk as root.
- Locate config file, which describes other partitions on the disk(s).
- Swap on (mkswap if necessary).
- Mount / (mkfs if necessary).
- Copy minimum system from ramdisk.
- pivot_root

Files on Disk

- System files (/bin, /lib, /etc, /sbin, /var) 28 Mb
- /local (/usr/local -> /local)
 - /local/home
- Symbolic links (/usr, /opt, many others)
- 521 files, 634 symbolic links

Files on Disk (booting)

- AFS file system starts.
- The 'package' command maintains directory structure.
- It also checks all resident files for size, date, permissions, and ownership.
- It also checks all directories, links, devices.
- It also deletes all local scratch and temporary files (not including cache files).

Adaptation of package

- Locally written utility (n lines of C).
- Config file + hostname (FQ) + IP address.
- Package's own macro processor mpp.
- Package.
- Reboot option

Cache

- Feature of AFS file system.
- Set to 500 Mb maximum.
- Persistent across reboots.
- File access -> check -> use cached copy
|-> get a new copy

Cache (booting)

- AFS file system already accessible.
- A special script (per cluster/machine) may run after the local system is checked.
- /sbin/init starts after this.

Full System

- Copy of normal distribution, modified to remove networking and some init scripts.
- On Network (in AFS) 2.3 Gb
- Read only – no access controls.
- AFS ACLs can be used for licensed software, if necessary.
- Users and authentication (various)
- User filestore (various)

Users and Authentication

- In central LDAP system: pam_ldap
- Special UNIX user names for courses created at boot time: pam_unix
- AFS users: pam_afs

Filestore

- Novell (mounted as NCPFS)
- SAN (mounted as SMBFS)
- Local UNIX filestore set up at boot time
- AFS for system files and a few users

How to Look After It

- Installation
- Administration
- Routine maintenance
- Bugs and problems

Why is it This Way?

- Constraints of management
- Constraints of resources

Evaluation

- Not rocket science
- Delays
 - at boot time
 - when using any program for the first time on that machine
- It works

Thousands of Linux Installations (and only one administrator)

A Linux cluster client
for the University of Manchester

A V Le Blanc
Manchester Computing
University of Manchester

LeBlanc@mcc.ac.uk